arXiv:1509.08371v1 [cond-mat.stat-mech] 24 Sep 2015

Using Bayes formula to estimate rates of rare events in transition path sampling simulations

Pierre Terrier,[1] Mihai-Cosmin Marinica,[1] and Manuel Athènes[1]

[1] *CEA, DEN, Service de Recherches de Métallurgie Physique, UPSay, F-91191 Gif-sur-Yvette, France*

Transition path sampling is a method for estimating the rates of rare events in molecular systems based on the gradual transformation of a path distribution containing a small fraction of reactive trajectories into a biased distribution in which these rare trajectories have become frequent. Then, a multistate reweighting scheme is implemented to postprocess data collected from the staged simulations. Herein, we show how Bayes formula allows to directly construct a biased sample containing an enhanced fraction of reactive trajectories and to concomitantly estimate the transition rate from this sample. The approach can remediate the convergence issues encountered in free energy perturbation or umbrella sampling simulations when the transformed distribution insufficiently overlaps with the reference distribution.

I. INTRODUCTION

The frequencies or rates of thermally activated events are crucial parameters that control atomic transport in condensed matter at equilibrium and the long-term evolution of many systems driven out of equilibrium. Not surprisingly, considerable effort has been devoted in the last decades to designing efficient molecular simulation methods for computing these rates. When the typical durations of the activated events are much shorter than the mean inter-event times, each event can be schematized as a transition of the system from a reactant basin (a) to a product basin (b) and the entire transiton paths can be studied using molecular dynamics (MD) simulations. So as to monitor a-to-b transitions in MD, the practitioner is first faced to the problem of defining the suitable functions taking input values in the configuration space Ω and indicating whether a particular configuration belongs to one of the two basins of interest. For this purpose, we usually consider the indicator function $h_{a(b)} : q \in \Omega \to \{0, 1\}$ whose output value is 1 if the input is in subset $a(b)$ and 0 elsewhere. This enables us to formalize an a-to-b time-correlation function as follows

$$C(t) \;=\; \frac{\int_0^{+\infty} h_a\left[q(s)\right] h_b\left[q(s+t)\right] ds}{\int_0^{+\infty} h_a\left[q(s)\right] ds} \qquad (1)$$

in which $q(s)$ is the system configuration at time s. When the fast molecular relaxation and the overall residence time within basin a occur on well-separated time scales, the derivative of the time-correlation function, $dC(t)/dt$, displays a transient plateau corresponding to the phenomenological rate for transitioning from a to b [1, 2]. The accurate estimation of $C(t)$ is a challenging task because transition events are rare on the simulation timescale. This problem can be alleviated by the use of the transition path sampling method [3, 4, 7–9] (TPS) in which a bias is introduced to enforce [3, 4] or favour [7–9] the sampling of the short trajectories transitioning from the product to the reactant basins. TPS writes the a-to-b time-correlation function as a conditional expectation of the indicator function of basin b at time t over an ensemble of short trajectories starting at equilibrium in basin a at time $t = 0$,

$$C(t) \;=\; \mathbb{E}\left\{h_b\left[q(t)\right] | h_a\left[q(0)\right] = 1\right\}. \qquad (2)$$

Let now write $\mathcal{O}$ the functional $h_b[q(t)]$ associated with paths starting in basin a at time $t = 0$. In principle, the time-correlation function in (2) can be estimated from the information contained in a set of M trajectories of duration t, after correcting for the simulation biases $\left\{e^{-B^m}\right\}_{1\leq m\leq M}$ using the standard reweighting scheme used in the umbrella sampling method [11, 12]

$$\widehat{\mathcal{O}}^M = \frac{\frac{1}{M}\sum_{m=1}^{M} \mathcal{O}^m e^{-B^m}}{\frac{1}{M}\sum_{m=1}^{M} e^{-B^m}}. \qquad (3)$$

The observation $\mathcal{O}^m$ is the value taken by the path observable $\mathcal{O}$ for the the mth trajectory. The standard reweighting scheme also relates to the free energy perturbation (FEP) method. [11, 12] The logarithm of the denominator in Eq. (3) corresponds to the free energy difference between the reference ensemble of interest and the biased (perturbed) ensemble that is sampled. To obtain an accurate estimate using (3), the biased sample must however contain typical data of the unbiased distribution in a significant proportion. Stated differently, the perturbed and unpertubed distributions should substantially overlap. If this condition is not met, then the associated free-energy difference is usually overestimated. [10–12]

In practice, the biased distributions that contains the reactive trajectories differs substantially from the reference distribution that contains non-reactive trajectories almost exclusively. As a result, the state-to-state correlation functions that have been estimated in TPS simulations so far were obtained through staged transformations. [4, 9] The biasing potential is gradually switched on and a simulation is performed at each stage of the switching protocol. Then, a postprocessing procedure (the weighted histogram analysis method [13] or the multistate Bennett acceptance ratio method [14]) is used to combine the data from the multiple simulations and, based on accurate estimates of the successive free-energy

differences, to eventually extract a reliable estimate of the desired observable [11, 12].

In this article, we show how Bayes formula can be used to (i) estimate any path observable expectation, (ii) adaptively construct a biasing potential whose associated distribution contains trajectories that are both reactive and non reactive (iii) estimate the state-to-state correlation function from a single sample owing to good overlapping properties of the biased distribution. The article is organized as follows. The Bayesian expectation in the rare event context is derived in Sec. II. We then compare the approach to standard reweighting in Sec. III. The migration of a vacancy in a crystal serves as an illustration of the approach in Sec. IV.

II. ESTIMATING PATH OBSERVABLES

A. Path ensembles and probability distributions

A path z consists of a sequence of $L+1$ states: $z = \{q_\ell, p_\ell\}_{0\leq\ell\leq L}$ where q_ℓ and p_ℓ are the ℓth positions and ℓth momenta of state $x_\ell = \{q_\ell, p_\ell\}$ at time $t = \ell\tau$ with τ being the timestep of the considered MD scheme. From the conditional probability to generate z using the MD scheme given initial state x_0, we define a path action by

$$\mathcal{H}(z) = H(x_0) - \beta^{-1}\ln \mathrm{P}_{\mathrm{MD}}(z|x_0). \tag{4}$$

where H is the system Hamiltonian and $\mathrm{P}_{\mathrm{MD}}(z|x_0)$ is the probability to generate path z knowing that the system at $t = 0$ is state x_0. In addition, a bias $\mathcal{K}(\theta, z)$, controlled by an external parameter θ, acts upon path z by modifying its occurrence probability. Confining the positions q_0 of the initial state of the path to basin a, the biased path distribution given θ writes

$$\pi(z|\theta) = h_a(q_0)\exp\left[\mathcal{A}(\theta) - \beta\mathcal{H}(z) - \mathcal{K}(\theta, z)\right] \tag{5}$$

where the function $\mathcal{A}(\theta)$ acts as a normalizing constant in path space, denoted by $\mathcal{Z}$:

$$\exp\left[-\mathcal{A}(\theta)\right] = \int_{\mathcal{Z}} \exp\left[-\beta\mathcal{H}(z) - \mathcal{K}(\theta, z)\right]\mathcal{D}z \tag{6}$$

In a first TPS protocol, [4] referred to as *confining protocol*, the bias aims at confining the trajectory endpoints into windows distributed along a reaction coordinate (RC), a function $\xi : \mathcal{Q} \to \mathbb{R}$ that is able to describe the transition pathway from basin a to basin b. Resorting to a family of indicator functions h^ξ_θ, the conditional probability writes $\pi(z|\theta) = h_a(q_0)h^\xi_\theta(q_L)\exp\left[\mathcal{A}(\theta) - \beta\mathcal{H}(z)\right]$, entailing

$$\mathcal{K}(\theta, z) \;=\; \begin{cases} 0 & \text{if } h^\xi_\theta(q_L) = 1, \\ +\infty & \text{otherwise.} \end{cases} \tag{7}$$

In Eq. (7), paths ending outside window h^ξ_θ have zero probability (5) and θ usually takes values in a finite integer set $\Theta = \{0, 1, 2, \cdots, \theta_{\max}\}$. Window h^ξ_0 ideally contains the whole phase space : $\forall q \in \mathcal{Q}$, the configuration space, we have $h^\xi_0(q) = 1$. Hence $\mathcal{K}(0, z)$ is always zero and $\pi(z|0)$ corresponds to the unbiased path probability distribution. On the other end, window associated with $\theta_{\max}$ maps basin b, i.e. $h^\xi_{\theta_{\max}} = h_b$. This implies that $\pi(z|\theta_{\max})$ is the probability distribution of the transition path ensemble and that the relation $C(L\tau) = \exp[-\mathcal{A}(\theta_{\max}) + \mathcal{A}(0)]$ holds.

In a second TPS protocol, referred to as *tilting protocol*, the bias acting upon the paths is the product of the external parameter and a path functional [6–9]

$$\mathcal{K}(\theta, z) = \theta\mathcal{L}(z). \tag{8}$$

Since trajectory endpoints are not constrained, the fraction of reactive paths with respect to $\pi(z|\theta_{\max})$ distribution is no more equal to one. For instance, in the recent set-up [9] that will be considered below in Sec. IV, $\mathcal{L}(z)$ is a non-positive functional that is all the more negative that the explored portions of the potential energy surface are more negatively curved. Because basins of attraction are separated from each others by mechanically unstable regions (corresponding to the negatively curved portions of the energy surface), the biasing approach is able to increase the occurrence of reactive trajectories up to 10−40%, depending on the choice of $\theta_{\max}$. [9] Although the equality between $C(L\tau)$ and $\exp[-\mathcal{A}(\theta_{\max}) + \mathcal{A}(0)]$ does not hold anymore, the free energy difference $\mathcal{A}(\theta_{\max}) - \mathcal{A}(0)$ is again to be determined to estimate the unbiased correlation function. [9]

Unfortunately, the associated free-energy difference cannot be estimated directly through free energy perturbation from a single sample of unbiased trajectories because the correlation function is very small, nor from a sample of biased trajectories because typical non-reactive paths are not generated. That is why TPS constructs additional samples of trajectories, confined using intermediate windows [4] or tilted using intermediate values for the external parameter. [9] It finally resorts to a rematching procedure to extract the $\mathcal{A}(\theta_{\max}) - \mathcal{A}(0)$. This computational bottleneck results from the insufficient overlap between the biased and unbiased distributions, as illustrated below in Sec. III within the tilting protocol (8).

We now show that it is practically possible to construct a biased distribution exhibiting adequate overlapping properties. Assuming that Θ is $[0, \theta_{\max}]$ interval, this probability distribution reads

$$\bar{\mathrm{P}}_{\mathcal{A}}(z) = \frac{1}{\theta_{\max}}\int_\Theta \pi(z|d\theta) \tag{9}$$

where the normalizing factor $1/\theta_{\max}$ corresponds to the uniform density of $\pi(z|\theta)$-distributions with respect to θ. Evaluating $\bar{\mathrm{P}}_{\mathcal{A}}(z)$ still requires the accurate knowledge of the free energy $\mathcal{A}$, since this quantity enters the sampled probability density (9) via (5). However, the overall approach is greatly simplified owing to the use of the adaptive biasing force (ABF) technique [16–19] within the method of expanded ensembles. [20, 21] The

latter approach introduces an auxiliary biasing potential denoted by $A(\theta)$, assumed to be constant first, and considers the extended space $\Theta \cup \mathcal{Z}$ equipped with path action $A(\theta) - \mathcal{K}(\theta, z)$. The occurrence of θ in the expanded ensemble defines the marginal probability of θ

$$\mathrm{P}_A(\theta) \propto \exp\left[A(\theta) - \mathcal{A}(\theta)\right]. \tag{10}$$

This relationship is obtained after integrating over the path space $\mathcal{Z}$ and plugging (6). It entails that, if A is strictly equal to the free energy, then the marginal probability of θ is constant and equal to $1/\theta_{\max}$ in Eq. (9). This last feature suggests a way of constructing $\bar{\mathrm{P}}_{\mathcal{A}}(z)$ since the ABF method permits to adapt $A(\theta)$ on $\mathcal{A}(\theta)$ via the current estimate of its derivative $\mathcal{A}'(\theta)$. Postponing the description of ABF technique to Sec. II C, we show how to estimate any observable expectation using Bayes formula. [22] This approach will require sampling the following marginal probability

$$\bar{\mathrm{P}}_A(z) = \int_\Theta \pi(z|\vartheta)\mathrm{P}_A(d\vartheta), \tag{11}$$

obtained after integrating $\vartheta \in \Theta$. This probability corresponds to the occurrence of any particular path in the expanded ensemble, irrespective of the value of θ. Using the marginal probabilities (10) and (11) and the conditional probability of z given θ, an analytical expression amenable to integration by numerical quadrature can be derived for the conditional probability of θ given z from Bayes formula

$$\bar{\pi}_A(\theta|z) \;=\; \frac{\pi(z|\theta)\mathrm{P}_A(\theta)}{\bar{\mathrm{P}}_A(z)} = \frac{e^{A(\theta)-\mathcal{K}(\theta,z)}}{\int_\Theta e^{A(\vartheta)-\mathcal{K}(\vartheta,z)}d\vartheta}. \tag{12}$$

Note that in (12) the conditional probability of z knowing θ does not depend on the auxiliary biasing potential [see expression Eq. (5)], while the three other involved probabilities do. We take advantage of this property and rearrange (12) so as to cast Bayes identity into a computationally useful form

$$\pi(z|\theta) \;=\; \frac{\bar{\pi}_A(\theta|z)\bar{\mathrm{P}}_A(z)}{\mathrm{P}_A(\theta)}. \tag{13}$$

This relation will enable one to express the conditional expectation given θ as an expectation involving the conditional probabilities of θ given the set of paths sampled according to probability (11).

B. Reweighting estimator based on Bayes formula

To show how this can be done, we first express the marginal probability of θ as an expectation involving the conditional probability of θ over the distribution of the marginal probability of z:

$$\pi(z|\theta) \;=\; \frac{\bar{\pi}_A(\theta|z)\bar{\mathrm{P}}_A(z)}{\int_{\mathcal{Z}} \bar{\pi}_A(\theta|\tilde{z})\bar{\mathrm{P}}_A(\mathcal{D}\tilde{z})} \tag{14}$$

This equation can be used to cast $\mathbb{E}_\pi\left(\mathcal{O}|\theta\right) = \int_{\mathcal{Z}} \mathcal{O}(z)\pi(\mathcal{D}z|\theta)$, the conditional expectation of any path observable $\mathcal{O}(z)$ given θ into the following form

$$\mathbb{E}_\pi\left(\mathcal{O}|\theta\right) \;=\; \frac{\int_{\mathcal{Z}} \mathcal{O}(z)\bar{\pi}_A(\theta|z)\bar{\mathrm{P}}_A(\mathcal{D}z)}{\int_{\mathcal{Z}} \bar{\pi}_A(\theta|z)\bar{\mathrm{P}}_A(\mathcal{D}z)} \tag{15}$$

Here, $\mathcal{O}(z)$ can be $h_b(q_\ell)$ the a-to-b time-correlation function or any path observable, such as $\mathcal{K}(\theta, z)$ or $\mathcal{L}(z)$ for instances.

Let now assume that a sample of trajectories $\{z^m\}_{1\le m\le M}$ is constructed using a Monte Carlo scheme obeying detailed balance in the path ensemble of probability density $\bar{\mathrm{P}}_A$. Applying the ergodic theorem to the ensemble average ratio (15), an estimator of the conditional expectation is

$$\hat{\mathbb{E}}_\pi^{\mathrm{B},M}\left(\mathcal{O}|\theta\right) = \frac{\frac{1}{M}\sum_{m=1}^M \mathcal{O}^m \bar{\pi}_A(\theta|z^m)}{\frac{1}{M}\sum_{m=1}^M \bar{\pi}_A(\theta|z^m)} \tag{16}$$

where $\mathcal{O}^m = \mathcal{O}(z^m)$ and superscript B stands for Bayes formula. This reweighting approach has been previously used in a molecular context to extract free energy profiles. [22] It was termed adiabatic reweighting in reference to the dynamical decoupling that is involved in molecular dynamics simulations between the external parameter and the particle coordinates (See Ref. 22).

C. Adaptive biasing force method

In ABF [16–19], the derivative of the auxiliary biasing potential with respect to θ is adapted on the current estimate of $\mathcal{A}'(\theta)$ the mean force along θ. In the long term, A converges to $\mathcal{A}$ (up to an additive constant) and $\mathrm{P}_A(\theta)$ becomes a uniform distribution. Differentiating $\mathcal{A}(\theta)$ is very simple mathematically within the tilting protocol (8). The mean force corresponds to the conditional expectation of $\mathcal{L}$ given θ

$$\mathcal{A}'(\theta) = \int_{\mathcal{Z}} \partial_\theta \mathcal{K}(\theta, z)\pi(\mathcal{D}z|\theta) = \mathbb{E}(\mathcal{L}|\theta). \tag{17}$$

In the application of Sec. IV, K replicas of the system will be simulated on a parallel computer architecture. We thus adapt the mean force $\mathcal{A}'(\theta)$ for $\theta \in \Theta$ following Ref. 22

$$\begin{aligned} A'_{\tilde{m}}(\theta) \;&=\; \frac{\sum_{k=1}^K \sum_{m=1}^{\tilde{m}-1} \mathcal{L}(z^{km})\pi_{A_m}(\theta|z^{km})}{\sum_{k=1}^K \sum_{m=1}^{\tilde{m}-1} \pi_{A_m}(\theta|z^{km})} \\ &=\; \frac{\sum_{k=1}^K \sum_{m=1}^{\tilde{m}-1} \mathcal{L}^{km} \frac{\exp\left[A_m(\theta)-\theta\mathcal{L}^{km}\right]}{\int_\Theta \exp[A_m(\vartheta)-\vartheta\mathcal{L}^{km}]d\vartheta}}{\sum_{k=1}^K \sum_{m=1}^{\tilde{m}-1} \frac{\exp[A_m(\theta)-\theta\mathcal{L}^{km}]}{\int_\Theta \exp[A_m(\vartheta)-\vartheta\mathcal{L}^{km}]d\vartheta}} \end{aligned} \tag{18}$$

The evaluation of the adaptive biasing force, a simple task compared to the evaluation of the interatomic forces, is shared by all cores. This is why the replica index k is absent in A'_m. The auxiliary potential is then integrated

for the next step. After the preliminary ABF run, the auxilary biasing potential A is frozen and the state-to-state correlation function is estimated in a subsequent run by setting $\mathcal{O}(z)$ to $h_b(q_\ell)$ $(0 \leq \ell \leq L)$ and θ to 0. In this situation, the generic BF estimator writes

$$\hat{\mathbb{E}}_\pi^{\mathrm{B},KM}(\mathcal{O}|\theta) = \frac{\frac{1}{KM}\sum_{k=1}^{K}\sum_{m=1}^{M}\mathcal{O}^{km}\bar{\pi}_A(\theta|z^{km})}{\frac{1}{KM}\sum_{k=1}^{K}\sum_{m=1}^{M}\bar{\pi}_A(\theta|z^{km})} \quad (19)$$

where z^{km} refers to the mth path of the kth replica and $\mathcal{O}^{km} = \mathcal{O}(z^{km})$.

Note that ABF is primarily used to compute mean forces along one- and two-dimensional RC, which requires to evaluate the first and second derivatives of the RC with respect to q. [17, 19] This task may be impractical in many circumstances, or even impossible when the RC is discrete for instance. Hence, it is not possible to implement ABF with the confining protocol (7) when the second derivative of ξ is not defined. This limitation however does not restrict the scope of TPS because trajectory endpoints can be softly confined to basin b by resorting to the tilting protocol (8) and setting $\mathcal{L}(z)$ to $\xi(q_\ell)$, the RC value of the path end-point. This feature is illustrated in next subsection.

D. Confining path endpoints through the tilting protocol

Let consider a generic model consisting of a particle evolving according to Brownian motion along a line (the momentum is omitted and the potential energy is constant). The position of the particle is obtained by integrating the corresponding overdamped Langevin equation: we have $q_{\ell+1} = q_\ell + \sqrt{2\tau D}B_\ell$ where τ is the timestep, D a diffusion coefficient and B_ℓ is a random variate drawn is the normal distribution of zero mean and unit variance. We set a to $\{0\}$ and and b to $[1, +\infty($. This implies that $q_0 = 0$ and q_L is distributed according to the normal distribution $q \to \sqrt{\omega/\pi}\exp(-\omega q^2)$ where $\omega = (4DL\tau)^{-1}$. Integrating this distribution from 1 to $+\infty$ then yields the probability that a path ends at a position larger than 1. This quantity defines the time-correlation $C(L\tau) = \mathrm{erfc}(\sqrt{\omega})/2$ that we wish to compute through the biased sampling of trajectory endpoints. When ω is large, diffusion is very slow and the probability $C(L\tau)$ to reach $q_L \geq 1$ is very small. We thus bias the simulation by adding the soft restraint $\mathcal{L}(z) = -2\omega q_L$, so as to gradually increase the fraction of trajectories ending in the product basin $q_L \geq 1$ with increasing θ value. This way of proceeding is similar in spirit to the confining protocol (7) which instead is based on a series of hard constraints.

In this illustration, we represent a path only using its endpoint, denoted by q to simplify. Plugging $\mathcal{A}(\theta) = \ln(\omega/\pi) - \omega\theta^2$ in Eq. (5), the conditional probability of q given θ is equal to

$$\pi(q|\theta) = \sqrt{\omega/\pi}\exp\left[-\omega(q-\theta)^2\right]. \quad (20)$$

The unbiased distribution is obtained by setting θ to 0 in (20) and is represented as a function of q by the green curve delimiting the red area in Fig. 1. The conditional distributions with various biases, obtained by setting θ to $\frac{1}{4}$, $\frac{1}{2}$, $\frac{3}{4}$ and 1 in (20) are displayed.

Two Brownian motions are considered: a fast one for which $\omega = 5$ and a slow one for which $\omega = 100$. The various distributions for the two values of the ω parameter are displayed in Fig. 1.a and in Fig. 1.b, respectively. For $\omega = 5$, we observe that the biased distributions $q \to \pi(q|1)$ substantially overlap with both the target region $(q \geq 1)$ and the reference distribution $q \to \pi(q|0)$ (areas displayed in red and green respectively). For $\omega = 100$, none of the biased distribution ($\theta = \frac{1}{4}$, $\frac{1}{2}$ and $\frac{3}{4}$) substantially overlaps with both the reference distribution and the target region. At variance, the marginal distribution (9) obtained for $\theta_{\max} = 1$ and both values of ω (curves displayed in blue) overlap with the unbiased distribution.

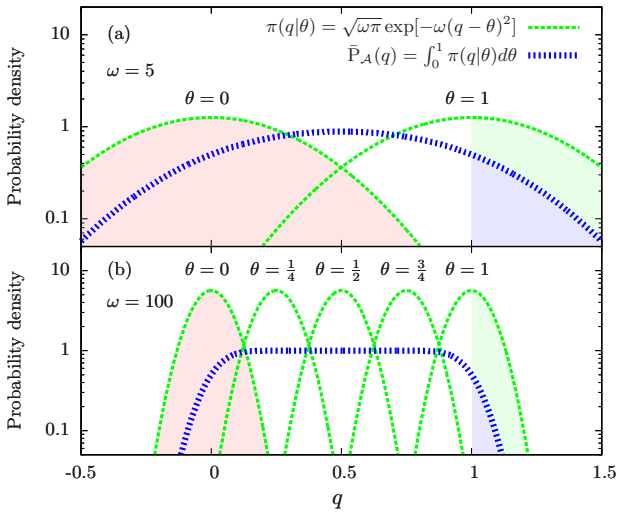

FIG. 1: Conditional distributions $\pi(q|\theta)$ for the indicated values of θ (green curves) and the marginal probability of q (blue curves) for $\omega = 5$ in pannel (a) and $\omega = 100$ in pannel (b). The target region corresponds to $q \geq 1$ and the unbiased distribution of q is filled in red.

A question then naturally arises as to how deviations from our choice of $\mathcal{A}$ as biasing potential qualitatively affects the overlapping properties of the marginal probability of q. To answer, we set the biasing potential to $(1+\delta)\mathcal{A}$: the deviation parameter δ determines how close the biasing potential is to the potential of mean force, the value zero of δ having been considered so far. We investigate in particular the effect of δ on the overlapping properties of the maginal probability distributions with respect to the reference distribution and the region of interest. We observe in Fig. 2 that the extent of overlap with the reference distribution increase with increasing values of δ and that the occurrence of an event concommitantly decreases. Hence, the choice $\delta = 0$ offers an interesting qualitative trade-off: the overlap with the ref-

erence distribution is substantial and the rare events are frequently observed.

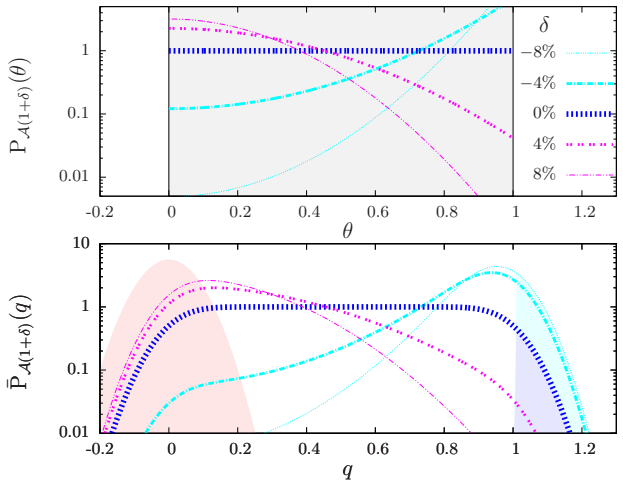

FIG. 2: Marginal probabilities of θ (a) and of q (b) for the biasing potentials $A = (1+\delta)\mathcal{A}$ with varying values of δ.

Next, the efficiency of the BF estimator is compared to the one of the corresponding standard estimator (SR). The latter one is derived in Sec. III. We use the present toy model [Sec. II D] for computing the a-to-b correlation function and investigate the relevant case where the auxiliary biasing potential is set to the potential of mean force. The efficiency of the ABF procedure II C for computing the potential of mean force is demonstrated in Sec. IV in simulations of vacancy diffusion.

III. ESTIMATOR EFFICIENCY

A. Standard reweighting estimator

The aforementioned method of expanded ensembles [20, 21] consists of sampling the extended space $(\vartheta, z) \in \Theta \cup \mathcal{Z}$ equipped with joint probability

$$\mathrm{p}_A(\vartheta, z) \propto h_a(q_0) \exp\left[A(\vartheta) - \beta\mathcal{H}(z) - \mathcal{K}(\vartheta, z)\right] \quad (21)$$

which is formally equal to $\pi(z|\vartheta)\mathrm{P}_A(\vartheta)$. This sampling approach alternatively suggests using the standard reweighting estimator in order to correct for the sampling biases. However, the SR estimator does not converge well in the present rare-event context, whether the biased sample be generated according to the joint probability distribution (21) or to the restricted distribution $\pi(z|\theta_{\max})$, wherein ϑ is set to $\theta_{\max}$. This limitation is illustrated in Sec. III B and III C by performing simulations in the expanded ensemble so as to allow direct comparison with the (more efficient) BF estimator presented in Sec. II B. The SR estimator associated with observable $\mathcal{O}$ is constructed from the following expectation ratio

$$\mathbb{E}_\pi(\mathcal{O}|\theta) = \frac{c_\Theta \int_{\Theta\cup\mathcal{Z}} \mathcal{O}(z) \frac{\mathrm{p}_A(\theta, z)}{\mathrm{p}_A(\vartheta, z)} \mathrm{p}_A(d\vartheta, \mathcal{D}z)}{c_\Theta \int_{\Theta\cup\mathcal{Z}} \frac{\mathrm{p}_A(\theta, z)}{\mathrm{p}_A(\vartheta, z)} \mathrm{p}_A(d\vartheta, \mathcal{D}z)}. \quad (22)$$

The normalizing constant $c_\Theta = 1/\left(\int_\Theta d\vartheta\right)$ is added to the numerator and denominator of ratio (22) so that the latter quantity is formally equal to $\mathrm{P}_A(\theta)$, the marginal probability of θ. Resorting to the ergodic theorem, the conditional expectation of observable $\mathcal{O}$ given θ and $\{\vartheta^m, z^m\}_{1\le m\le M}$, a Markov chain constructed using a Monte Carlo scheme leaving the joint probability distribution (21) invariant, may be estimated using (SR estimator)

$$\hat{\mathbb{E}}_\pi^{\mathrm{S},M}(\mathcal{O}|\theta) = \frac{\frac{c_\Theta}{M} \sum_{m=1}^{M} \mathcal{O}^m e^{A(\theta)-\mathcal{K}(\theta,z^m)} \Big/ e^{A(\vartheta^m)-\mathcal{K}(\vartheta^m,z^m)}}{\frac{c_\Theta}{M} \sum_{m=1}^{M} e^{A(\theta)-\mathcal{K}(\theta,z^m)} \Big/ e^{A(\vartheta^m)-\mathcal{K}(\vartheta^m,z^m)}}. \quad (23)$$

Remarkably, the BF estimator (23) can also be implemented when the sample is generated according to the joint distribution. To justify its use, one simply resorts to the ergodic theorem within the following expectation ratio:

$$\mathbb{E}_\pi(\mathcal{O}|\theta) = \frac{\int_{\Theta\cup\mathcal{Z}} \mathcal{O}(z)\bar{\pi}_A(\theta|z)\mathrm{p}_A(d\vartheta, \mathcal{D}z)}{\int_{\Theta\cup\mathcal{Z}} \bar{\pi}_A(\theta|z)\mathrm{p}_A(d\vartheta, \mathcal{D}z)} \quad (24)$$

Relation (24) is a consequence of Bayes formula as it is obtained by plugging $\bar{\mathrm{P}}_A(z) = \int_\Theta \mathrm{p}_A(d\vartheta, z)$ into (15).

Estimates associated with the unbiased path distribution are recovered by setting θ to 0 in (16). The BF estimator simplifies to

$$\widehat{\mathcal{O}}^{\mathrm{B},M} = \frac{\frac{1}{M} \sum_{m=1}^{M} \mathcal{O}^m e^{A(0)} \Big/ \left(\int_\Theta e^{A(\vartheta)-\mathcal{K}(\vartheta,z^m)} d\vartheta\right)}{\frac{1}{M} \sum_{m=1}^{M} e^{A(0)} \Big/ \left(\int_\Theta e^{A(\vartheta)-\mathcal{K}(\vartheta,z^m)} d\vartheta\right)} \quad (25)$$

Resorting to the sequence of biasing potentials $B^m = \ln \int_\Theta \exp\left[A(\vartheta) - \mathcal{K}(\vartheta, z^m)\right] d\vartheta$, we may write estimator in the conventional form (3) mentioned in the introduction. The SR estimator with respect to the unbiased path distribution ($\theta = 0$) reads

$$\widehat{\mathcal{O}}^{\mathrm{S},M} = \frac{\frac{c_\Theta}{M} \sum_{m=1}^{M} \mathcal{O}^m e^{A(0)-A(\vartheta^m)+\mathcal{K}(\vartheta^m,z^m)}}{\frac{c_\Theta}{M} \sum_{m=1}^{M} e^{A(0)-A(\vartheta^m)+\mathcal{K}(\vartheta^m,z^m)}}. \quad (26)$$

The denominators in Eqn. (25) and (26) both yield estimates of $\mathrm{P}_A(\theta)$ with $\theta = 0$. Relevant information about the estimator convergence can be learned by monitoring the marginal probability of θ.

Note that $A(\vartheta^m) - \mathcal{K}(\vartheta^m, \mathcal{L}^m)$ corresponds to the biasing potential associated with path z^m. Compared with

the SR estimator (26), the sampled values of the external parameter θ are irrelevant information in the BF estimator (25). We now show that these distinct features results in different numerical efficiencies.

B. Comparison of BF and SR estimators in a generic model

We consider the toy model of Sec. II D in which the biasing potential has already converged to the potential of mean force and $\theta_{\max} = 1$, implying that the marginal probability density of θ is equal to one in Θ and that the conditional probability of q given θ in (20) is equal to the joint probability $\pi(q|\theta) = \sqrt{\omega/\pi}\exp\left(-\omega(q-\theta)^2\right)$.

The correlation function $C(L\tau)$ is estimated using BF estimator (25) and SR estimator (26) by setting the observable $\mathcal{O}$ to $h_b(q)$ equal to 1 if $q \geq 1$ and to 0 otherwise. The marginal probability of θ at 0 is also estimated from the denominator of (25) and (26). Each generated Markov chain $\{\vartheta^m, q^m\}_{1\leq m\leq M}$ is used twice, first to obtain a BF estimate based on (25) and then to get a SR estimate based on (26). The mth state of the chain is generated as follows: ϑ^m is drawn randomly and uniformly in $\Theta = [0,1]$ interval and q^m is drawn in the Gaussian distribution of $(2\omega)^{-1}$ variance and ϑ^m mean. Displayed in Fig. (3) are the means and standard errors of 10^5 independent estimates, obtained using both estimators. We observe that, with increasing ω parameter, only the BF estimator yields an accurate estimation of the marginal probability of θ at 0 (Fig. 3.a) and of the correlation function (3.b). The computational speed-up of convergence that is achieved by using BF estimator rather than SR estimator can be assessed from their respective standard errors plotted as a function of ω in Fig. 3.c. As soon as ω becomes larger than 20, the standard error associated with BF estimator is two orders of magnitude lower than the one obtained using the SR estimator. The use of Bayes formula accelerates the simulations by about four orders of magnitude.

C. Fluctuation relations

To explain why BF and SR estimators behave differently, we resort to fluctuation relations [10, 23] in order to analyse how the quantity $\mathrm{P}_{\mathcal{A}}(0)$ is evaluated in both approaches. This probability corresponds to the denominator of ratio (25) for the BF estimator, and of ratio (26) for the SR estimator, when A is set to $\mathcal{A}$. The probability density is one when Θ is the $[0,1]$ interval. This probability is computed from two different ensemble averages, entailing two distinct types of fluctuation relations. For the BF approach, the following single relation must hold

$$\langle e^{-\Delta(q)}\rangle = \int_{\mathbb{R}} e^{-\Delta}\bar{\Pi}(d\Delta) = e^0 \qquad (27)$$

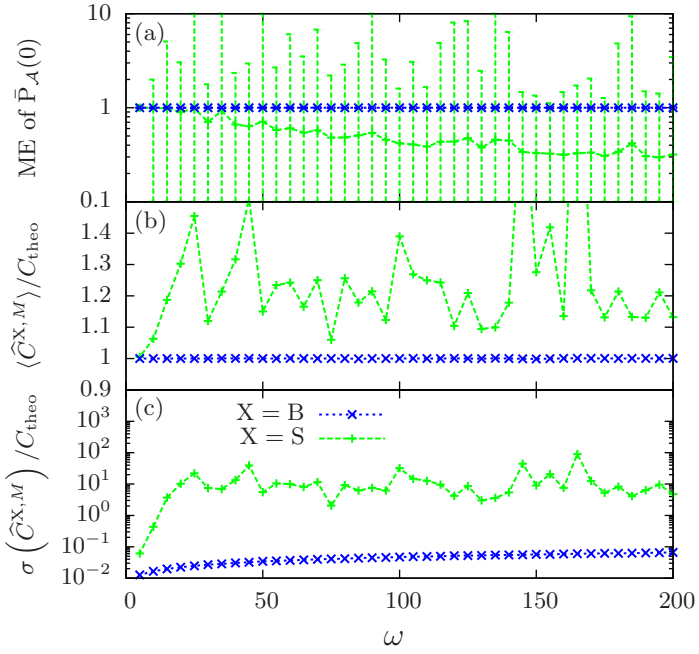

FIG. 3: Comparison between the Bayesian (16) and standard (26) estimators: (a) averaged estimates of the marginal probability of θ at 0 as a function of ω, (b) averaged estimates of the probability $C(M\tau)$ normalized to the exact probability and (c) its normalized standard error. All averages are obtained from 10^5 estimates and each estimate is obtained using $M = 10^4$ points.

where $\Delta(q) = -\ln\left[\mathrm{P}_{\mathcal{A}}(0,q)/\bar{\mathrm{P}}_{\mathcal{A}}(q)\right]$ and $\bar{\Pi}(\Delta)$ denotes the probability that $\Delta(q)$ takes the particular value Δ. As for the SR approach, a whole family of detailed fluctuation relations must be satisfied by the quantity $\Delta(\vartheta,q) = -\ln\left[\mathrm{P}_{\mathcal{A}}(0,q)/\mathrm{P}_{\mathcal{A}}(\vartheta,q)\right]$

$$\langle e^{-\Delta(\theta,q)}\rangle_\theta = \int_{\mathbb{R}} e^{-\Delta}\Pi_\theta(d\Delta) = e^0, \qquad (28)$$

$$\langle e^{-\Delta(\vartheta,q)}\rangle = \int_{\Theta}\langle e^{-\Delta(\theta,q)}\rangle_\theta \mathrm{P}_{\mathcal{A}}(d\theta) = e^0 \qquad (29)$$

where $\Pi_\theta(\Delta)$ denotes the probability that $\Delta(\theta,q)$ takes the particular value Δ. The uniform average over $\theta\in\Theta$ yields an additional overall fluctuation relation (29) corresponding to (27). Because the exponential function is strictly increasing, the negative values of Δ and the stricly positive ones have similar statistical weight in (27) or (28), in the sense that the sum of the two contibutions are equal to one (except for the case $\theta = 0$). Hence, excessively small fractions of negative Δ values will result in large statistical variance and in slow numerical convergence of the estimates as a function of simulation time. [10] The probabilities $\Pi_\theta(\Delta)$ with $\theta\in\left\{0,\frac{1}{4},\frac{1}{2},\frac{3}{4},1\right\}$ and $\bar{\Pi}(\Delta)$ have been plotted as a function of Δ for $\omega = 5$ in Fig. 4.a and and for $\omega = 100$ in Fig. 4.b. From the distributions at $\omega = 100$, we observe that the probability to have $\Delta(\vartheta^m, q^m) \leq 0$ is negligible when $\vartheta^m > \frac{1}{2}$. Given the fact that the ϑ^m's are sampled uniformly in $[0,1]$,

substantial deviations from fluctuation relations (28) will inevitably be measured in typical (finite-length) simulations, resulting in inaccurate SR estimates. At variance, the BF approach does not suffer from this limitation, as the fluctuation relation that must be satisfied is global. We indeed observe that the fraction of the negative values of Δ is always substantial, making the BF estimator particularly efficient for large ω values.

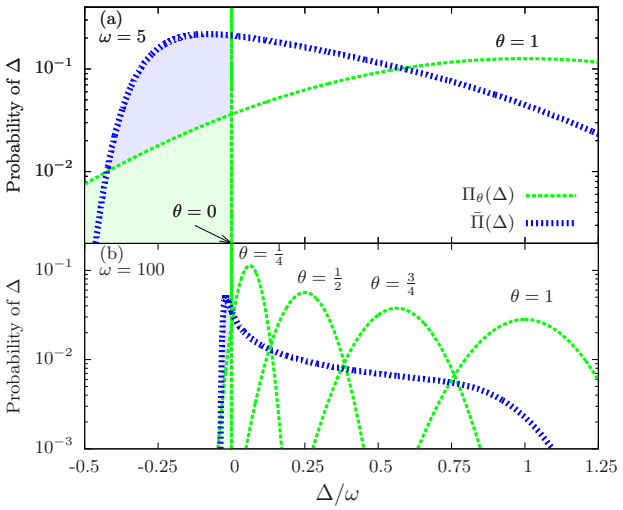

FIG. 4: Probability distribution of Δ as a function of the Δ, blue curve for BF estimator and green curves for SR estimator with various values of θ.

IV. APPLICATION TO VACANCY MIGRATION IN α-FE

We now demonstrate the efficiency of the approach in simulations of the migration of a single vacancy on a lattice in α-Fe, a crystalline phase of iron with body-centered cubic (BCC) structure. The migration corresponds to the jump of an atom in the [111] direction into a nearest neighbor vacant site. The length of the jump is $a_0\sqrt{3}/2$, where a_0 is the side of the BCC cube. Atomic interactions of this atomic system are described by an embedded atom model potential. [24] Reference values for migration rates are available for this testbed system that was previously investigated in Ref. 9 using a combination of transition path sampling and the multistate Bennett acceptance ratio method for postprocessing.

A. Computational set-up

The computational set-up is as follows. Basin a and b are defined with respect to the underlying perfect lattice whose sites are the atomic positions of the structure at 0 K without the vacancy. The lattice parameter is $a_0 = 2.8553$Å. The indicator function h_a is equal to 1 if all atoms are located within a distance of 0.45Å from their lattice site, and to 1 otherwise. The characteristic function h_b is 1 if one atom is located beyond a distance of $a_0\sqrt{3}/4$ from its lattice site, otherwise it is 0. Path sampling consists of shooting and shifting moves as detailed in the Appendices A and B. Trajectories contain $L = 150$ steps with time-step $\tau = 2fs$. A position-Verlet scheme [9] is used to construct $x_{\ell+1}$ from x_ℓ, meaning that the gradient of the potential energy is evaluated at $q_{\ell+1/2} = q_\ell + p_\ell\tau/2$. The Jacobian matrix associated with the MD transformation exhibits eigenvalues that are either complex numbers located on the unit circle or real positive numbers. Let us denote by $\mu_{\ell+1/2}$ the logarithm of the smallest eigenvalue modulus. Its value is characterized by the eigenvalue spectrum of the Hessian matrix associated with the potential energy at $q_{\ell+1/2}$. Details on the connection between the Hessian and Jacobian matrices are given in Ref. 9. The value of $\mu_{\ell+1/2}$ is strictly negative when the lowest egeinvalue of the Hessian matrix is strictly negative, in which case the energy surface is negatively curved along the direction generated by the corresponding eigenvector. The biasing path functional is set to

$$\mathcal{L}(z) = \max\left[\mathcal{L}_{\min}, \sum_{\ell=0}^{L-1} \mu_{\ell+1/2}\right] \leq 0 \qquad (30)$$

The cut-off parameter $\mathcal{L}_{\min}$ is set to the value -9. It is used to prevent from exploring regions containing second order saddles and thus to save computational time, trajectories leading to such regions corresponding to non-reactive rare events. The lowest eigenvalues of the Hessian is computed using the Lanczos algorithm. [26] Details about the numerical implementation are given in Ref. 27, 28.

The illustrations are given at the temperature of 500 K. We set $\theta_{\max} = 2.1$. In the following, a simulation run utilizes $K = 480$ replicas and consists of $M = 10^4$ Monte Carlo cycles. Each replica is allocated to a distinct processor. A cycle consists of performing a shooting move followed by a shifting move for each replica. The two procedures are detailed in Appendices A and B, as their implementation slightly differs from the one given previoulsy [9].

B. Construction of the auxiliary potential

Two consecutive series of 5 independent simulation runs are performed. The first five runs aim at constructing the auxiliary biasing potential using the ABF scheme (18). Then, freezing the previoulsy obtained biasing potentials, 5 subsequent (production) runs are performed to estimate the expectations using the scheme (19). Figure 5 displays the estimates of $P_{\mathcal{A}}(\theta)$, $\mathcal{A}(\theta)$ and $\mathbb{E}(\mathcal{L}|\theta)$ as a function of θ and averaged over the 5 runs. The standard errors are evaluated from the 5 estimates and are indicated by error bars for the three quantities in Fig. 5. We observe that a flat histogram

is obtained for the marginal probability of θ. Reproductible data are obtained for the mean force and its potential. Furthermore, the difference between the adaptive and production runs is insignificant, final averages could have been taken after the adaptation run.

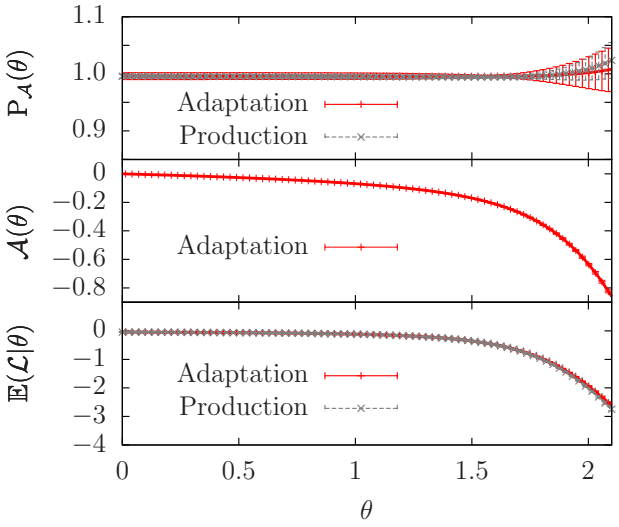

FIG. 5: Marginal probability of θ (multiplied by $\theta_{\max}$), potential of mean force and mean force as a function of θ.

The standard errors are small and not clearly visible on the graphs in Fig. 5, except on the curve displaying the marginal probability of θ. We observe that the standard errors associated with $P_{\mathcal{A}}(\theta)$ increase with θ and become substantial at large θ values. To explain this trend, let us examine $\bar{\rho}(\mathcal{L})$, the probability distributions of the sampled $\mathcal{L}$ values. We observe in Fig. 6 that this distribution is bimodal. The large peak at 0 corresponds to typical trajectories that are non reactive. The smaller peak in the range from -7 to -4 contains both reactive trajectories and active trajectories returning to a. The presence of two peaks means that metastability is not completely suppressed through path-sampling, even though the fraction of reactive trajectories is enhanced by several orders of magnitude compared with the one associated with the unbiased distribution $\rho(\mathcal{L}|0)$. Concerning the biased distribution $\rho(\mathcal{L}|\theta_{\max})$, the peak containing the reactive trajectories is higher and more pronounced than that of the sampled distribution. This feature explaining the substantial statistical fluctuations observed in the measurement of $P_{\mathcal{A}}(\theta)$ when θ is large. Note that the θ_c value for which the two peaks of the bimodal distribution $\rho(\mathcal{L}|\theta_c)$ have equal weights occurs in the range $2.2 - 2.4$ and decreases with the path length [9]. Here, θ_c would correspond to the inflexion of the $\mathbb{E}(\mathcal{L}|\theta)$ curve, outside the plot in Fig. 5.c. As reported in Ref. 9, the restricted sampling of the conditional distribution $\pi(z|\theta)$ becomes very difficult when $\theta > \theta_c$, the measured autocorrelation function of $\mathcal{L}$ increases drastically. Important autocorrelations are also observed in the sampling of $\bar{P}_{\mathcal{A}}(z)$ distribution when $\theta_{\max}$ is set to a value larger than θ_c. However, the extent of metastability is smaller with the approach based on Bayes formula. We speculate that this results from the smaller barrier height for trajectory disactivation for $\bar{\rho}(\mathcal{L})$ distribution than for $rho(\mathcal{L}|\theta)$ distribution, as indicated in Fig. 6 by the red and blue downward arrows, respectively.

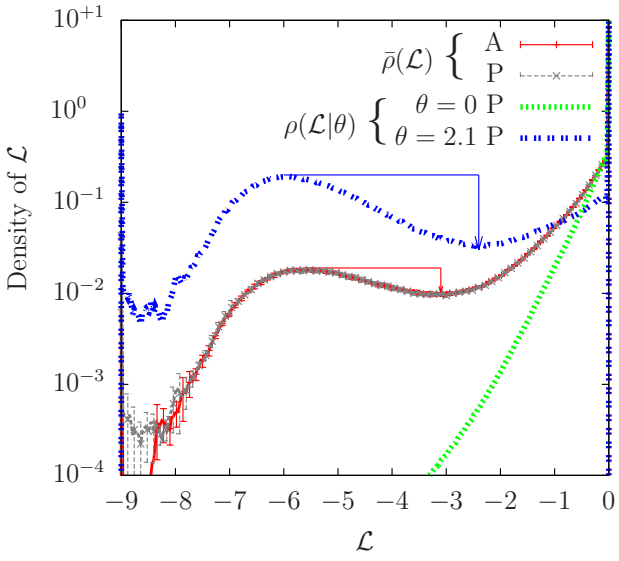

FIG. 6: Distributions of $\mathcal{L}$: $\bar{\rho}(\mathcal{L})$ and $\rho(\mathcal{L}|\theta)$ denote the probabilities that $\mathcal{L}(z)$ takes value $\mathcal{L}$ with respect to path distributions $\bar{P}_{\mathcal{A}}(z)$ and $\pi(z|\theta)$, respectively.

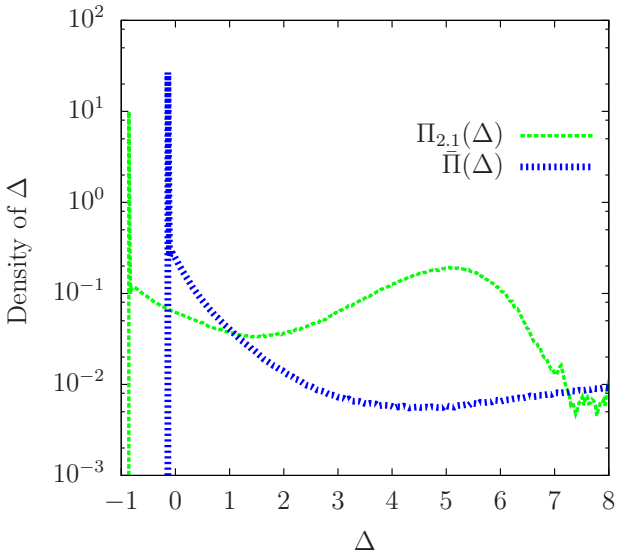

FIG. 7: Distribution of Δ defined in Sec. III C.

Another advantage of the approach based on Bayes formula is illustrated from the fluctuation theorems defined in Sec. III C. We observe in Fig. 7 that the distributions of the Δ values are more peaked around 0 when the marginal probability is sampled, a result expected and consistent with the previous analysis using the toy model. Numerical convergence should therefore be faster for the reason mentioned in Sec. III

The sampled path distributions with $\theta_{\max} < \theta_c$ contain high enough a fraction of reactive trajectories so as to accurately estimate the a-to-b correlation function. Figure 8 represents the time correlation function and its time derivative as obtained after the 5 adaptation runs and the 5 production runs. Smooth values are obtained for the time-derivative owing to recycling of the shifted trajectories using the waste-recycling procedure described in Appendix C. The phenomenological transition rate corresponds to the plateau value, which is in perfect agreement with the value previousy calculated in Ref. 9.

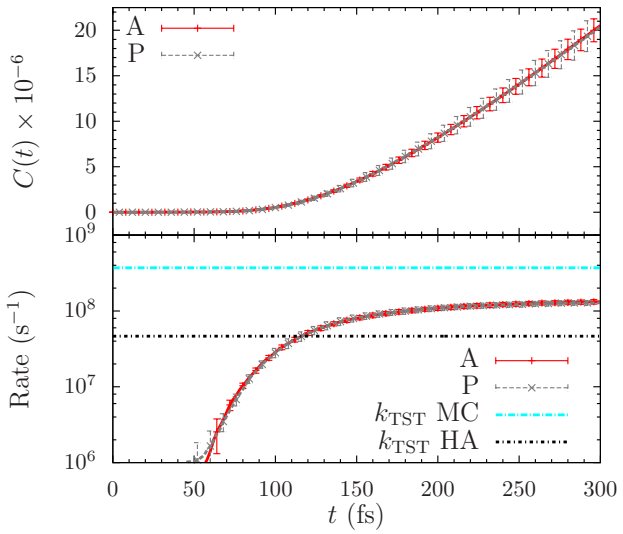

FIG. 8: Time correlation function and its derivative as a function of the time $t = \ell\tau$, where ℓ is the index of the trajectory states. The horizontal lines corresponds to the values obtained using transition state theory [9, 29] in which the free energy barrier is evaluated using either Monte Carlo simulations (MC) or the classical harmonic approximation (HA).

The present simulations involving the hopping of a vacancy in a crystal show that the approach can easily be implemented on a parallel computer architecture by propagating several replicas of the system simultaneously and adapting the biasing force periodically. The proposed framework allows to explore the multiple reaction channels corresponding to the 8 possible atomic jumps into the vacancy. This is a clear advantage compared with other rare event approaches such as transition interface sampling, [30] and forward flux sampling [31] that tends to confine trajectories into separate transition channels. The use of Bayes formula in TPS method thus facilitates its implementation and will certainly extend its scope.

V. CONCLUSION

Estimating the state-to-state correlation functions in TPS method requires to extract the free energy difference associated with the transformation of a trajectory ensemble into a perturbed ensemble wherein the rare reactive trajectories have become frequent. This task had been achieved so far through postprocessing of the information contained in a series of Markov chains, in which the strength of the perturbating bias favouring the occurrence of reactive trajectories was gradually increased. TPS was therefore a most relevant case study to apply the recently proposed approach [22] for computing free energies based on Bayes formula and adaptive biasing. The combination of the two techniques allows to adaptively construct a biased sample whose associated distribution substantially overlaps with both the unpertubed and perturbed distributions and to *in fine* obtain an unbiased estimate of the time-correlation function.

The ability of the proposed approach to construct and sample a biased distribution that substantially overlaps with any pertubed and unperturbed distributions is extremely useful in general. From a larger perspective, this feature will enable one to automatically remediate the convergence issue that is encountered in umbrella sampling when the perturbed distribution, sampled by this FEP-based simulation method, insufficiently overlaps with the reference distribution. This last situation is quite ubiquitous in molecular simulation and, not incidentally, motivated the testbed model investigation of Sec. III. We thus expect many interesting applications of the Bayes formula approach in order to compute rare-event frequencies or mean-force potentials in domains ranging from chemistry to bio-physics and materials science.

Acknowledgments

Stimulating discussions with G. Stoltz are acknowledged. This work was performed using HPC resources from GENCI-[CCRT/CINES] (Grant x2015096973).

[1] D. Chandler, J. Chem. Phys. **68**, 2959 (1978).
[2] D. Chandler, *Introduction to modern statistical mechanics*, Oxford Univ. Press (1987).
[3] C. Dellago, P. G. Bolhuis and D. Chandler, J. Chem. Phys. **110**, 6617 (1999).
[4] P. G. Bolhuis, D. Chandler, C. Dellago and P. L. Geissler, Annu. Rev. Phys. Chem. **53**, 291 (2002).
[5] M. Athènes and V. V. Bulatov, Phys. Rev. Lett. **113**, 230601 (2014).
[6] M. Athènes, Eur. Phys. J. B **38**, 651 (2004).
[7] Y. Elmatad, R. Jack, D. Chandler and J. Garrahan, Proc. Natl. Acad. Sci. U.S.A. **107**, 12793 (2010).
[8] P. Geiger and C. Dellago, Chem. Phys. **375**, 309 (2010).
[9] M. Athènes, M.-C. Marinica and T. Jourdan J. Chem.

Phys. **137**, 194107 (2012).
[10] C. H. Bennett, J. Comp. Phys. **22**, 245 (1976).
[11] D. Frenkel and B. Smit, *Understanding molecular simulation: from algorithms to applications* Academic Press (2002).
[12] T. Lelièvre, M. Rousset and G. Stoltz, *Free-energy computations: a mathematical perspective*, Imperial College Press, 2010.
[13] A. M. Ferrenberg and R. H. Swendsen, Phys. Rev. Lett. **63**, 1195 (1989).
[14] M. J. Shirts and J. D. Chodera, J. Chem. Phys. **129**, 124105 (2008).
[15] G. Torrie and J. Valleau, J. Comput. Phys. **23**, 187 (1977).
[16] E. Darve and A. Pohorille, J. Chem. Phys. **115**, 9169 (2001).
[17] E. Darve, D. Rodriguez-Gomez, A. Pohorille, J. Chem. Phys. **128**, 144120 (2008).
[18] T. Lelièvre, M. Rousset, G. Stoltz, Nonlinearity **21**, 1155 (2008).
[19] J. Hénin, J. Fiorin, C. Chipot, M. Klein, J. Chem. Theory Comput. **6**, 35 (2010).
[20] A. Lyubartsev, A. Martinovskii, S. Shevkunov and P. Vorontsov-Velyaminov, J. Chem. Phys. **96**, 1776 (1992).
[21] Y. Iba, Int. J. Mod. Phys. C, **12**, 623 (2001).
[22] L. Cao, G. Stoltz, T. Lelièvre, M.-C. Marinica and M. Athènes, J. Chem. Phys. 140, 104108 (2014).
[23] M. Athènes, Eur. Phys. J. B **58**, 83 (2007).
[24] G. Ackland, M. I. Medelev, D. J. Srolovitz, S. Han and A. V. Barashev, J. Phys.: Condens. Matter **16**, 2629 (2004).
[25] M. Athènes and M.-C. Marinica, J. Comput. Phys. **229**, 7129 (2010).
[26] C. Lanczos, *Applied Analysis*, Prentice Hall (1961).
[27] E. Cances, F. Legoll, M.-C. Marinica, K. Minoukadeh and F. Willaime, J. Chem. Phys. **130**, 114711 (2009).
[28] M.-C. Marinica, F. Willaime and N. Mousseau, Phys. Rev. B, **83**, 094119 (2011).
[29] P. Hänggi, P. Talkner and M. Borkovec, Rev. Mod. Phys. **62**, 251 (1990).
[30] T. van Erp, D. Moroni and P. Bolhuis, J. Chem. Phys. **118**, 7762 (2003).
[31] R. Allen D. Frenkel and P. ten Wolde, J. Chem. Phys. **124**, 024102 (2006) & 194111 **124** (2006).
[32] G. Stoltz, J. Comp. Phys. **225**, 491 (2007).
[33] D. Frenkel, Proc. Natl. Acad. Sci. U.S.A. **101**, 17571 (2004).
[34] D. Frenkel, Waste-recycling Monte Carlo, in "Computer Simulations in Condensed Matter Systems", Lect. Notes Phys. **703**, 127 (2006).
[35] J.-F. Delmas and B. Jourdain, J. Appl. Probab. **46**, 938 (2009).
[36] M. Athènes and F. Calvo, Chem. Phys. Chem. **16**, 2332 (2008).
[37] G. Adjanor, M. Athènes, J. M. Rodgers, J. Chem. Phys. **135**, 044127 (2011).
[38] J. Kim, J. M. Rodgers, M. Athnes, and B. Smit, J. Chem. Theory Comput. **7**, 3208 (2011).

Appendix A: Shooting moves

A shooting move consists in performing the following operations (z is current path)

i. draw an integer ℓ randomly and uniformly in $\{0, 1, \cdots, L\}$

ii. from $x_\ell \in z$, generate a neighboring state $\tilde{x}_\ell$ according to the canonical distribution ;

iii. construct the trial trajectory $\tilde{z}$ by applying the Verlet map $L-\ell$ times forward from $\tilde{x}_\ell$ and ℓ time backward from $\tilde{x}_\ell$;

iv. compute the new path biasing potential

$$B(\tilde{z}) = \ln \int_\Theta \exp\left[A(\theta) - \mathcal{K}(\theta, \tilde{z})\right] d\theta \tag{A1}$$

v. draw a random number R uniformly in $(0, 1]$; if R is lower than the acceptance probability

$$\mathrm{P}_{\mathrm{acc}}[\tilde{z} \leftarrow z] = \min\left\{1, h_a(\tilde{q}_0) \exp\left[B(\tilde{z}) - B(z)\right]\right\}, \tag{A2}$$

accept the trial trajectory, otherwise reject it.

This scheme is correct because the probability to accept a trial path in (A2) corresponds to a proper Metropolis-hasting acceptance probability. We have

$$\mathrm{P}_{\mathrm{acc}}[\tilde{z} \leftarrow z] = \min\left\{1, \frac{\mathrm{P}_{\mathrm{gen}}[z \leftarrow \tilde{z}]\bar{\mathrm{P}}_A(\tilde{z})}{\mathrm{P}_{\mathrm{gen}}[\tilde{z} \leftarrow z]\bar{\mathrm{P}}_A(z)}\right\}, \tag{A3}$$

where $\mathrm{P}_{\mathrm{gen}}[\tilde{z} \leftarrow z]$ is the probability to generate $\tilde{z}$ from z and vice versa for $\mathrm{P}_{\mathrm{gen}}[z \leftarrow \tilde{z}]$. In practice, the path generating procedure in ii constructs new momenta $\tilde{p}_\ell$ in the canonical distribution by slightly perturbating the momenta p_ℓ using an Ornstein-Uhlenbeck process [9, 32]. We use $\tilde{p}_\ell = \epsilon p_\ell + \sqrt{1-\epsilon^2}\varpi$ where ϖ are momenta drawn in the Maxwell-Boltmzann distribution at temperature β^{-1}. The parameter $\epsilon \in [0, 1]$ mixes the uncorrelated momenta p_ℓ and ϖ. As a result, the following condition is satisfied

$$\frac{\mathrm{P}_{\mathrm{gen}}[z \leftarrow \tilde{z}]}{\mathrm{P}_{\mathrm{gen}}[\tilde{z} \leftarrow z]} = \frac{\exp[-\beta\mathcal{H}(\tilde{z})]}{\exp[-\beta\mathcal{H}(z)]}. \tag{A4}$$

Noticing that the marginal probabilities of path z and $\tilde{z}$ in (A3) are proportional to $h_a(q_0)e^{B(z)-\beta\mathcal{H}(z)}$ and $h_a(\tilde{q}_0)e^{B(\tilde{z})-\beta\mathcal{H}(\tilde{z})}$ respectively, and plugging (A4) into the formal rate (A3), the simple form (A2) is obtained. It is then easy to show that detailed balance with respect to the sampled distribution is obeyed. [11, 12] In the numerical application given in Sec. IV, the value of the mixing parameter is $\epsilon = 0.975$, which yields a mean acceptance rate of 90.6%.

Appendix B: Shifting moves

Let $\bar{z} = \{\bar{x}_h\}_{0 \le h \le 2L}$ denote a path of $2L$ states generated by the Verlet map and $\tilde{z}_\ell = \{\bar{x}_{h+\ell}\}_{0 \le h \le L}$ with $0 \le \ell \le L$ denote the $L+1$ path segments of L states included in $\bar{z}$. The path space for the extended paths is $\bar{\mathcal{Z}}$.

We define $\bar{\mathrm{Q}}(\bar{z}|z)$ the conditional probability to generate $\bar{z}$ from z using the Verlet map. We have

$$\bar{\mathrm{Q}}(\bar{z}|z) = \begin{cases} \frac{1}{L+1} & \text{if } \exists \ell \in \{0,1,\cdots,L\} : \tilde{z}_\ell = z \\ 0 & \text{otherwise.} \end{cases} \quad \text{(B1)}$$

The form of $\mathrm{Q}(\bar{z}|z)$ reflects the fact that there are exactly $1+L$ distinct paths $\bar{z} \in \bar{\mathcal{Z}}$ (generated through $2L$ successive applications of the Verlet map) such that $\bar{z}$ includes path $z \in \mathcal{Z}$. The marginal probability of $\bar{z}$ then follows

$$\bar{\mathcal{P}}_A(\bar{z}) = \int_{\mathcal{Z}} \bar{\mathrm{Q}}(\bar{z}|z)\bar{\mathrm{P}}_A(z)\mathcal{D}z \quad \text{(B2)}$$

$$= \frac{1}{L+1}\sum_{\ell=0}^{L} \bar{\mathrm{P}}_A(\tilde{z}_\ell) \quad \text{(B3)}$$

where $z \in \mathcal{Z}$. We also define the probability to select $\bar{z}_\ell$ from $\bar{z}$ in the shifting procedure as

$$\mathrm{P}_{\mathrm{sel}}(\ell|\bar{z}) = \frac{\mathrm{Q}(\bar{z}|\tilde{z}_\ell)\bar{\mathrm{P}}_A(\tilde{z}_\ell)}{\bar{\mathcal{P}}_A(\bar{z})} = \frac{h_a(\bar{x}_\ell)\exp\left[B(\tilde{z}_\ell)\right]}{\sum_{l=0}^{L} h_a(\bar{x}_l)\exp\left[B(\tilde{z}_l)\right]} \quad \text{(B4)}$$

A shifting move then consists in performing the following operations (z is current path)

i. draw an integer ℓ randomly and uniformly in $\{0,1,\cdots,L\}$;

ii. set $\bar{x}_h = x_{h-\ell}$ for $\ell \leq h \leq L+\ell$;

iii. construct $\{\bar{x}_h\}_{L+\ell<h\leq 2L}$ by applying the Verlet map $L-\ell$ times forward from $\bar{x}_{L+\ell}$;

iv. construct $\{\bar{x}_h\}_{0\leq h<\ell}$ by applying the Verlet map ℓ times backard from $\bar{x}_\ell$;

v. set $\bar{z} = \{\bar{x}_h\}_{0\leq h\leq 2L}$;

vi. select path $\tilde{z}_\sigma$ by drawing σ in the multinomial distribution of probability vector $\sigma \to \mathrm{P}_{\mathrm{sel}}(\sigma|\bar{z})$;

vii. update the state indexes of new path $z = \{x_h\}_{0\leq h\leq L}$ by setting $x_h = \bar{x}_{\sigma+h}$ for $0 \leq h \leq L$;

The proof that shifting moves leave the probability distribution $\bar{\mathrm{P}}_A(z)$ invariant follows the same lines as those given in Ref. 9. It can additionally be shown that the marginal probability $\bar{\mathcal{P}}_A(\tilde{z})$ is also left invariant by the samping procedures. This property is useful as it allows to estimate the observable using a waste-recycling estimator detailed hereafter.

Appendix C: Waste-recycling estimator

Using the shifting proposal probability defined in Appendix B, we express the conditional probability of (θ, z) given $\bar{z}$ using the following Bayes formula

$$\mu_A(\theta, z|\bar{z}) = \frac{\mathrm{Q}(\bar{z}|z)\mathrm{p}_A(\theta, z)}{\bar{\mathcal{P}}_A(\bar{z})}, \quad \text{(C1)}$$

$$= \pi_A(\theta|z) \times \frac{\mathrm{Q}(\bar{z}|z)\bar{\mathrm{P}}_A(z)}{\bar{\mathcal{P}}_A(\bar{z})}. \quad \text{(C2)}$$

When z is equal to $\tilde{z}_\ell$ ($\tilde{z}_\ell \in \bar{z}$), the conditional probability simplifies into

$$\mu_A(\theta, \ell|\bar{z}) = \pi_A(\theta|\tilde{z}_\ell)\mathrm{P}_{\mathrm{sel}}(\ell|\bar{z}). \quad \text{(C3)}$$

Denoting $\mathcal{O}(\tilde{z}_\ell^{km})$ by $\mathcal{O}_\ell^{km}$, the waste-recycling estimator at $\theta = 0$ is based on Bayes formula (C1) and writes

$$\widehat{\mathcal{O}}^{\mathrm{W},K,M} = \frac{\sum_{k=1}^{K}\sum_{m=1}^{M}\sum_{\ell=0}^{L}\mathcal{O}_\ell^{km}\mu_A(0,\ell|\tilde{z}^{km})}{\sum_{k=1}^{K}\sum_{m=1}^{M}\sum_{\ell=0}^{L}\mu_A(0,\ell|\tilde{z}^{km})}. \quad \text{(C4)}$$

The waste-recycling estimator can be used to estimate conditional expectations given θ in general. It is based on Bayes formula (C1) cast in its expectation form:

$$\mathbb{E}_\mu(\mathcal{O}|\theta) = \frac{\int_{\bar{\mathcal{Z}}}\sum_{\ell=0}^{L}\mathcal{O}(\tilde{z}_\ell)\mu(\theta,\ell|\bar{z})\bar{\mathcal{P}}_A(\mathcal{D}\bar{z})}{\int_{\bar{\mathcal{Z}}}\sum_{\ell=0}^{L}\mu(\theta,\ell|\bar{z})\bar{\mathcal{P}}_A(\mathcal{D}\bar{z})}. \quad \text{(C5)}$$

The fact that this expectation is equivalent to the expectation of interest can be checked by plugging

$$\sum_{\ell=0}^{L}\mathcal{O}(\tilde{z}_\ell)\mu(\theta,\ell|\bar{z}) = \int_{\mathcal{Z}}\mathcal{O}(z)\mu(\theta,z|\bar{z})\mathcal{D}z \quad \text{(C6)}$$

in Eq. (C5) and simplifying:

$$\mathbb{E}_\mu(\mathcal{O}|\theta) = \frac{\int_{\bar{\mathcal{Z}},\mathcal{Z}}\mathcal{O}(z)\mu(\theta,z|\bar{z})\bar{\mathcal{P}}_A(\bar{z})\mathcal{D}\bar{z}\mathcal{D}z}{\int_{\bar{\mathcal{Z}},\mathcal{Z}}\mu(\theta,z|\bar{z})\bar{\mathcal{P}}_A(\bar{z})\mathcal{D}\bar{z}\mathcal{D}z}$$

$$= \frac{\int_{\mathcal{Z}}\mathcal{O}(z)\mathrm{p}_A(\theta,\mathcal{D}z)}{\int_{\mathcal{Z}}\mathrm{p}_A(\theta,\mathcal{D}z)} \quad \text{(C7)}$$

$$= \int_{\mathcal{Z}}\mathcal{O}(z)\pi(\mathcal{D}z|\theta) = \mathbb{E}_\pi(\mathcal{O}|\theta). \quad \text{(C8)}$$

This procedure is adapted from [9]. For additional references on waste-recycling, see the original articles [33, 34], articles discussing connections with Bayes formulae [23, 25], a mathematical analysis of convergence [35] and some applications [36–38].